Non-asymptotic Analysis of ℓ_1-norm Support Vector Machines

Anton Kolleck, Jan Vybíral

arXiv:1509.08083v1 [cs.IT] 27 Sep 2015

Abstract

Support Vector Machines (SVM) with ℓ_1 penalty became a standard tool in analysis of highdimensional classification problems with sparsity constraints in many applications including bioinformatics and signal processing. Although SVM have been studied intensively in the literature, this paper has to our knowledge first non-asymptotic results on the performance of ℓ_1-SVM in identification of sparse classifiers. We show that a d-dimensional s-sparse classification vector can be (with high probability) well approximated from only $O(s\log(d))$ Gaussian trials. The methods used in the proof include concentration of measure and probability in Banach spaces.

Index Terms

Support vector machines, compressed sensing, machine learning, regression analysis, signal reconstruction, classification algorithms, functional analysis, random variables

I. Introduction

A. Support Vector Machines

Support vector machines (SVM) are a group of popular classification methods in machine learning. Their input is a set of data points $x_1,\dots,x_m\in\mathbb{R}^d$, each equipped with a label $y_i\in\{-1,+1\}$, which assigns each of the data points to one of two groups. SVM aims for binary linear classification based on separating hyperplane between the two groups of training data, choosing a hyperplane with separating gap as large as possible.

Since their introduction by Vapnik and Chervonenkis [27], the subject of SVM was studied intensively. We will concentrate on the so-called soft margin SVM [8], which allow also for misclassification of the training data are the most used version of SVM nowadays.

In its most common form (and neglecting the bias term), the soft-margin SVM is a convex optimization program

$$\min_{\substack{w\in\mathbb{R}^d\\ \xi\in\mathbb{R}^m}} \frac{1}{2}\|w\|_2^2+\lambda\sum_{i=1}^m\xi_i \quad \text{subject to} \quad y_i\langle x_i,w\rangle\geq 1-\xi_i \quad \text{and} \quad \xi_i\geq 0 \tag{I.1}$$

for some tradeoff parameter $\lambda>0$ and so called slack variables ξ_i. It will be more convenient for us to work with the following equivalent reformulation of (I.1)

$$\min_{w\in\mathbb{R}^d}\sum_{i=1}^m[1-y_i\langle x_i,w\rangle]_+ \quad \text{subject to} \quad \|w\|_2\leq R, \tag{I.2}$$

where $R>0$ gives the restriction on the size of w. We refer to monographs [25], [28], [29] and references therein for more details on SVM and to [13, Chapter B.5] and [9, Chapter 9] for a detailed discussion on dual formulations.

B. ℓ_1-SVM

As the classical SVM (I.1) and (I.2) do not use any pre-knowledge about w, one typically needs to have more training data than the underlying dimension of the problem, i.e. $m\gg d$. Especially in analysis of high-dimensional data, this is usually not realistic and we typically deal with much less training data, i.e. with $m\ll d$. On the other hand, we can often assume some structural assumptions on w, in the most simple case that it is *sparse*, i.e. that most of its coordinates are zero. Motivated by the success of LASSO [26] in sparse linear regression, it was proposed in [6] that replacing the ℓ_2-norm $\|w\|_2$ in (I.2) by its ℓ_1-norm $\|w\|_1=\sum_{j=1}^d|w_j|$ leads to sparse classifiers $w\in\mathbb{R}^d$. This method was further popularized in [34] by Zhu, Rosset, Hastie, and Tibshirani, who developed an algorithm that efficiently computes the whole solution path (i.e. the solutions of (I.2)

A. Kolleck is with the Department of Mathematics, Technical University Berlin, Street of 17. June 136, 10623 Berlin, Germany, (e-mail:kolleck@math.tu-berlin.de). A. Kolleck was supported by the DFG Research Center Matheon "Mathematics for key technologies" in Berlin.

J. Vybíral is with the Department of Mathematical Analysis, Charles University, Sokolovská 83, 186 00, Prague 8, Czech Republic, (e-mail: vybiral@karlin.mff.cuni.cz). J. Vybíral was supported by the ERC CZ grant LL1203 of the Czech Ministry of Education and by the Neuron Fund for Support of Science.

for a wide range of parameters $R > 0$). We refer also to [5], [2], [18] and [19] for other generalizations of the concept of SVM.

Using the ideas of concentration of measure [20] and random constructions in Banach spaces [21], the performance of LASSO was analyzed in the recent area of compressed sensing [11], [7], [3], [10], [12].

ℓ_1-SVM (and its variants) found numerous applications in high-dimensional data analysis, most notably in bioinformatics for gene selection and microarray classification [30], [31], [15]. Finally, ℓ_1-SVM's are closely related to other popular methods of data analysis, like elastic nets [32] or sparse principal components analysis [33].

C. Main results

The main aim of this paper is to analyze the performance of ℓ_1-SVM in the non-asymptotic regime. To be more specific, let us assume that the data points $x_1, \dots, x_m \in \mathbb{R}^d$ can be separated by a hyperplane according to the given labels $y_1, \dots, y_m \in \{-1, +1\}$, and that this hyperplane is normal to a s-sparse vector $a \in \mathbb{R}^d$. Hence, $\langle a, x_i \rangle > 0$ if $y_i = 1$ and $\langle a, x_i \rangle < 0$ if $y_i = -1$. We then obtain $\hat{a}$ as the minimizer of the ℓ_1-SVM. The first main result of this paper (Theorem II.3) then shows that $\hat{a}/\|\hat{a}\|_2$ is a good approximation of a, if the data points are i.i.d. Gaussian vectors and the number of measurements scales linearly in s and logarithmically in d.

Later on, we introduce a modification of ℓ_1-SVM by adding an additional ℓ_2-constraint. It will be shown in Theorem IV.1, that it still approximates the sparse classifiers with the number of measurements m growing linearly in s and logarithmically in d, but the dependence on other parameters improves. In this sense, this modification outperforms the classical ℓ_1-SVM.

D. Organization

The paper is organized as follows. Section II recalls the concept of ℓ_1-Support Vector Machines of [34]. It includes the main result, namely Theorem II.3. It shows that the ℓ_1-SVM allows to approximate sparse classifier a, where the number of measurements only increases logarithmically in the dimension d as it is typical for several reconstruction algorithms from the field of compressed sensing. The two most important ingredients of its proof, Theorems II.1 and II.2, are also discussed in this part. The proof techniques used are based on the recent work of Plan and Vershynin [24], which in turn makes heavy use of classical ideas from the areas of concentration of measure and probability estimates in Banach spaces [20], [21].

Section III gives the proofs of Theorems II.1 and II.2. In Section IV we discuss several extensions of our work, including a modification of ℓ_1-SVM, which combines the ℓ_1 and ℓ_2 penalty.

Finally, in Section V we show numerical tests to demonstrate the convergence results of Section II. In particular, we compare different versions of SVM and 1-Bit Compressed Sensing, which was first introduced by Boufounos and Baraniuk in [4] and then discussed and continued in [23], [24], [22], [1], [17] and others.

E. Notation

We denote by $[\lambda]_+ := \max(\lambda, 0)$ the positive part of a real number $\lambda \in \mathbb{R}$. By $\|w\|_1, \|w\|_2$ and $\|w\|_\infty$ we denote the ℓ_1, ℓ_2 and ℓ_∞ norm of $w \in \mathbb{R}^d$, respectively. We denote by $\mathcal{N}(\mu, \sigma^2)$ the normal (Gaussian) distribution with mean μ and variance σ^2. When ω_1 and ω_2 are random variables, we write $\omega_1 \sim \omega_2$ if they are equidistributed. Multivariate normal distribution is denoted by $\mathcal{N}(\mu, \Sigma)$, where $\mu \in \mathbb{R}^d$ is its mean and $\Sigma \in \mathbb{R}^{d \times d}$ is its covariance matrix. By $\log(x)$ we denote the natural logarithm of $x \in (0, \infty)$ with basis e. Further notation will be fixed in Section II under the name of "Standing assumptions", once we fix the setting of our paper.

II. ℓ_1-NORM SUPPORT VECTOR MACHINES

In this section we give the setting of our study and the main results. Let us assume that the data points $x_1, \dots, x_m \in \mathbb{R}^d$ are equipped with labels $y_i \in \{-1, +1\}$ in such a way that the groups $\{x_i : y_i = 1\}$ and $\{x_i : y_i = -1\}$ can indeed be separated by a sparse classifier a, i.e. that

$$y_i = \operatorname{sign}(\langle x_i, a \rangle), \quad i = 1, \dots, m \tag{II.1}$$

and

$$\|a\|_0 = \#\{j : a_j \neq 0\} \leq s. \tag{II.2}$$

As the classifier is usually not unique, we cannot identify a exactly by any method whatsoever. Hence we are interested in a good approximation of a obtained by ℓ_1-norm SVM from a minimal number of training data. To achieve this goal, we will assume that the training points

$$x_i = r\tilde{x}_i, \quad \tilde{x}_i \sim \mathcal{N}(0, \mathrm{Id}) \tag{II.3}$$

are i.i.d. measurement vectors for some constant $r > 0$.

To allow for more generality, we replace (II.2) by

$$\|a\|_2 = 1, \quad \|a\|_1 \leq R. \tag{II.4}$$

Let us observe, that $\|a\|_2 = 1$ and $\|a\|_0 \leq s$ implies also $\|a\|_1 \leq \sqrt{s}$, i.e. (II.4) with $R = \sqrt{s}$.

Furthermore, we denote by $\hat{a}$ the minimizer of

$$\min_{w \in \mathbb{R}^d} \sum_{i=1}^{m} [1 - y_i \langle x_i, w \rangle]_+ \quad \text{subject to} \quad \|w\|_1 \leq R. \tag{II.5}$$

Let us summarize the setting of our work, which we will later on refer to as "Standing assumptions" and which we will keep for the rest of this paper.

Standing assumptions:

(i) $a \in \mathbb{R}^d$ is the true (nearly) sparse classifier with $\|a\|_2 = 1, \quad \|a\|_1 \leq R$, $R \geq 1$, which we want to approximate;
(ii) $x_i = r\tilde{x}_i, \quad \tilde{x}_i \sim \mathcal{N}(0, \mathrm{Id}), i = 1, \ldots, m$ are i.i.d. training data points for some constant $r > 0$;
(iii) $y_i = \mathrm{sign}(\langle x_i, a \rangle), \quad i = 1, \ldots, m$ are the labels of the data points;
(iv) $\hat{a}$ is the minimizer of (II.5);
(v) Furthermore, we denote

$$K = \{w \in \mathbb{R}^d \mid \|w\|_1 \leq R\}, \tag{II.6}$$

$$f_a(w) = \frac{1}{m} \sum_{i=1}^{m} [1 - y_i \langle x_i, w \rangle]_+, \tag{II.7}$$

where the subindex a denotes the dependency of f_a on a (via y_i).

In order to estimate the difference between a and $\hat{a}$ we adapt the ideas of [24]. First we observe

$$\begin{aligned} 0 &\leq f_a(a) - f_a(\hat{a}) \\ &= \big(\mathbb{E}f_a(a) - \mathbb{E}f_a(\hat{a})\big) + \big(f_a(a) - \mathbb{E}f_a(a)\big) \\ &\qquad - \big(f_a(\hat{a}) - \mathbb{E}f_a(\hat{a})\big) \\ &\leq \mathbb{E}(f_a(a) - f_a(\hat{a})) + 2 \sup_{w \in K} |f_a(w) - \mathbb{E}f_a(w)|, \end{aligned}$$

i.e.

$$\mathbb{E}(f_a(\hat{a}) - f_a(a)) \leq 2 \sup_{w \in K} |f_a(w) - \mathbb{E}f_a(w)|. \tag{II.8}$$

Hence, it remains

- to bound the right hand side of (II.8) from above and
- to estimate the left hand side in (II.8) by the distance between a and $\hat{a}$ from below.

We obtain the following two theorems, whose proofs are given in Section III.

Theorem II.1. *Let $u > 0$. Under the "Standing assumptions" it holds*

$$\sup_{w \in K} |f_a(w) - \mathbb{E}f_a(w)| \leq \frac{8\sqrt{8\pi} + 18rR\sqrt{2\log(2d)}}{\sqrt{m}} + u$$

with probability at least

$$1 - 8\left(\exp\left(\frac{-mu^2}{32}\right) + \exp\left(\frac{-mu^2}{32r^2R^2}\right)\right).$$

Theorem II.2. *Let the "Standing assumptions" be fulfilled and let $w \in K$. Put*

$$c = \langle a, w \rangle, \quad c' = \sqrt{\|w\|_2^2 - \langle a, w \rangle^2}$$

and assume that $c' > 0$. If furthermore $c \leq 0$, then $\pi\mathbb{E}(f_a(w) - f_a(a))$ can be estimated from below by

$$\frac{\pi}{2} + c'r\frac{\sqrt{\pi}}{\sqrt{2}} - \frac{\sqrt{2\pi}}{r}.$$

If $c > 0$, then $\pi\mathbb{E}(f_a(w) - f_a(a))$ can be estimated from below by

$$\frac{\sqrt{\pi}}{\sqrt{2}} \int_0^{1/cr} (1 - crt) e^{\frac{-t^2}{2}}\, dt + \frac{c'}{c} \exp\left(\frac{-1}{2c^2r^2}\right) - \frac{\sqrt{2\pi}}{r}.$$

Combining Theorems II.1 and II.2 with (II.8) we obtain our main result.

Theorem II.3. *Let $d \geq 2$, $0 < \varepsilon < 0.18$, $r > \sqrt{2\pi}(0.57 - \pi\varepsilon)^{-1}$ and $m \geq C\varepsilon^{-2}r^2R^2\log(d)$ for some constant C. Under the "Standing assumptions" it holds*

$$\frac{\left\| a - \frac{\hat{a}}{\|\hat{a}\|_2} \right\|_2}{\langle a, \frac{\hat{a}}{\|\hat{a}\|_2} \rangle} \leq C' \left(\varepsilon + \frac{1}{r} \right) \tag{II.9}$$

with probability at least

$$1 - \gamma \exp\left(-C'' \log(d)\right) \tag{II.10}$$

for some positive constants γ, C', C''.

Remark II.4. 1) If the classifier $a \in \mathbb{R}^d$ with $\|a\|_2 = 1$ is s-sparse, we always have $\|a\|_1 \leq \sqrt{s}$ and we can choose $R = \sqrt{s}$ in Theorem II.3. The dependence of m, the number of samples needed, is then linear in s and logarithmic in d. Intuitively, this is the best what we can hope for. On the other hand, we leave it open, if the dependence on ε and r is optimal in Theorem II.3.

2) Theorem II.3 uses the constants C, C' and C'' only for simplicity. More explicitly we show that taking

$$m \geq 4\varepsilon^{-2} \left(8\sqrt{8\pi} + 19rR\sqrt{2\log(2d)} \right)^2,$$

we get the estimate

$$\frac{\|a - \hat{a}/\|\hat{a}\|_2\|_2}{\langle a, \hat{a}/\|\hat{a}\|_2 \rangle} \leq 2e^{1/2} \left(\pi\varepsilon + \frac{\sqrt{2\pi}}{r} \right)$$

with probability at least

$$1 - 8\bigg(\exp\bigg(\frac{-r^2R^2\log(2d)}{16} \bigg) + \exp\bigg(\frac{-\log(2d)}{16} \bigg) \bigg).$$

3) If we introduce an additional parameter $t > 0$ and choose $m \geq 4\varepsilon^{-2}(8\sqrt{8\pi} + (18+t)rR\sqrt{2\log(2d)})^2$, nothing but the probability changes to

$$1 - 8\bigg(\exp\bigg(\frac{-t^2r^2R^2\log(2d)}{16} \bigg) + \exp\bigg(\frac{-t^2\log(2d)}{16} \bigg) \bigg).$$

Hence, by fixing t large, we can increase the value of C'' and speed up the convergence of (II.10) to 1.

Proof of Theorem II.3: To apply Theorem II.1 we choose

$$u = \frac{rR\sqrt{2\log(2d)}}{\sqrt{m}}$$

and

$$m \geq 4\varepsilon^{-2}(8\sqrt{8\pi} + 19rR\sqrt{2\log(2d)})^2$$

and we obtain the estimate

$$\sup_{w \in K} |f_a(w) - \mathbb{E}f_a(w)| \leq \frac{8\sqrt{8\pi} + 18rR\sqrt{2\log(2d)}}{\sqrt{m}} + u \leq \frac{\varepsilon}{2}$$

with probability at least

$$\begin{aligned} &1 - 8\bigg(\exp\bigg(\frac{-mu^2}{32} \bigg) + \exp\bigg(\frac{-mu^2}{32r^2R^2} \bigg) \bigg) \\ &= 1 - 8\bigg(\exp\bigg(\frac{-r^2R^2\log(2d)}{16} \bigg) + \exp\bigg(\frac{-\log(2d)}{16} \bigg) \bigg). \end{aligned}$$

Using (II.8) this already implies

$$\mathbb{E}\big(f_a(\hat{a}) - f_a(a) \big) \leq \varepsilon \tag{II.11}$$

with at least the same probability. Now we want to apply Theorem II.2 with $w = \hat{a}$ to estimate the left hand side of this inequality. Therefore we first have to deal with the case $c' = \sqrt{\|\hat{a}\|_2^2 - \langle a, \hat{a} \rangle^2} = 0$, which only holds if $\hat{a} = \lambda a$ for some

$\lambda \in \mathbb{R}$. If $\lambda > 0$, then $\hat{a}/\|\hat{a}\|_2 = a$ and the statement of the Theorem holds trivially. If $\lambda \leq 0$, then the condition $f(\hat{a}) \leq f(a)$ can be rewritten as

$$\sum_{i=1}^{m}[1 + |\lambda| \cdot |\langle x_i, a\rangle|]_+ \leq \sum_{i=1}^{m}[1 - |\langle x_i, a\rangle|]_+.$$

This inequality holds if, and only if, $\langle x_i, a\rangle = 0$ for all $i = 1, \ldots, m$ - and this in turn happens only with probability zero.

We may therefore assume that $c' \neq 0$ holds almost surely and we can apply Theorem II.2. Here we distinguish the three cases $c = \langle \hat{a}, a\rangle \leq 0$, $0 < c \leq 1/r$ and $1/r < c$. First, we will show that the two cases $c \leq 0$ and $0 < c < 1/r$ lead to a contradiction and then, for the case $c > 1/r$, we will prove our claim.

1. case $c \leq 0$*:* Using Theorem II.2 we get the estimate

$$\pi\mathbb{E}(f_a(\hat{a}) - f_a(a)) \geq \frac{\pi}{2} + c'r\frac{\sqrt{\pi}}{\sqrt{2}} - \frac{\sqrt{2\pi}}{r} \geq \frac{\pi}{2} - \frac{\sqrt{2\pi}}{r}$$

and (II.11) gives (with our choices for r and ε) the contradiction

$$\frac{1}{\pi}\left(\frac{\pi}{2} - \frac{\sqrt{2\pi}}{r}\right) \leq \mathbb{E}(f_a(\hat{a}) - f_a(a)) \leq \varepsilon.$$

2. case $0 < c \leq 1/r$*:* As in the first case we use Theorem II.2 in order to show a contradiction. First we get the estimate

$$\begin{aligned}
&\pi\mathbb{E}(f_a(\hat{a}) - f_a(a)) \\
&\geq \frac{\sqrt{\pi}}{\sqrt{2}}\int_0^{1/cr}(1 - crt)e^{\frac{-t^2}{2}}\,dt + \frac{c'}{c}\exp\left(\frac{-1}{2c^2r^2}\right) - \frac{\sqrt{2\pi}}{r} \\
&\geq \frac{\sqrt{\pi}}{\sqrt{2}}\int_0^{1/cr}(1 - crt)e^{\frac{-t^2}{2}}\,dt - \frac{\sqrt{2\pi}}{r}.
\end{aligned}$$

Now we consider the function

$$g\colon (0,\infty) \to \mathbb{R}, \quad z \mapsto \int_0^{1/z}(1 - zt)e^{\frac{-t^2}{2}}\,dt.$$

It holds $g(z) \geq 0$ and

$$g'(z) = -\int_0^{1/z} te^{\frac{-t^2}{2}}\,dt < 0,$$

so g is monotonic decreasing. With $cr < 1$ this yields

$$\begin{aligned}
\pi\mathbb{E}(f_a(\hat{a}) - f_a(a)) &\geq \frac{\sqrt{\pi}}{\sqrt{2}}\int_0^{1/cr}(1 - crt)e^{\frac{-t^2}{2}}\,dt - \frac{\sqrt{2\pi}}{r} \\
&= \frac{\sqrt{\pi}}{\sqrt{2}}g(cr) - \frac{\sqrt{2\pi}}{r} \geq \frac{\sqrt{\pi}}{\sqrt{2}}g(1) - \frac{\sqrt{2\pi}}{r} \\
&= \frac{\sqrt{\pi}}{\sqrt{2}}\int_0^{1}(1 - t)e^{\frac{-t^2}{2}}\,dt - \frac{\sqrt{2\pi}}{r} \\
&\geq 0.57 - \frac{\sqrt{2\pi}}{r}.
\end{aligned}$$

Again, (II.11) now gives the contradiction

$$\frac{1}{\pi}\left(0.57 - \frac{\sqrt{2\pi}}{r}\right) \leq \mathbb{E}(f_a(\hat{a}) - f_a(a)) \leq \varepsilon.$$

We conclude that it must hold $c' > 0$ and $c > 1/r$ almost surely.

3. case $1/r < c$*:* In this case we get the estimate

$$\begin{aligned}
\pi\mathbb{E}(f_a(\hat{a}) - f_a(a)) &\geq \frac{\sqrt{\pi}}{\sqrt{2}}\int_0^{1/cr}(1 - crt)e^{\frac{-t^2}{2}}\,dt \\
&\qquad + \frac{c'}{c}\exp\left(\frac{-1}{2c^2r^2}\right) - \frac{\sqrt{2\pi}}{r} \\
&\geq \frac{c'}{c}\exp\left(\frac{-1}{2c^2r^2}\right) - \frac{\sqrt{2\pi}}{r} \\
&\geq \frac{c'}{c}e^{-1/2} - \frac{\sqrt{2\pi}}{r},
\end{aligned} \tag{II.12}$$

where we used $cr > 1$ for the last inequality. Further we get

$$\begin{aligned}\frac{c'}{c} &= \frac{\sqrt{\|\hat{a}\|_2^2 - \langle a, \hat{a}\rangle^2}}{\langle a, \hat{a}\rangle} = \sqrt{\frac{\|\hat{a}\|_2^2 - \langle a, \hat{a}\rangle^2}{\langle a, \hat{a}\rangle^2}} \\ &= \sqrt{\left(\frac{\|\hat{a}\|_2 - \langle a, \hat{a}\rangle}{\langle a, \hat{a}\rangle}\right)\left(\frac{\|\hat{a}\|_2 + \langle a, \hat{a}\rangle}{\langle a, \hat{a}\rangle}\right)} \\ &= \sqrt{\frac{(2 - 2\langle a, \hat{a}/\|\hat{a}\|_2\rangle)(2 + 2\langle a, \hat{a}/\|\hat{a}\|_2\rangle)}{4\langle a, \hat{a}/\|\hat{a}\|_2\rangle^2}} \\ &= \sqrt{\frac{\|a - \hat{a}/\|\hat{a}\|_2\|_2^2 \cdot \|a + \hat{a}/\|\hat{a}\|_2\|_2^2}{4\langle a, \hat{a}/\|\hat{a}\|_2\rangle^2}} \\ &\geq \frac{1}{2}\frac{\|a - \hat{a}/\|\hat{a}\|_2\|_2}{\langle a, \hat{a}/\|\hat{a}\|_2\rangle}.\end{aligned} \tag{II.13}$$

Finally, combining (II.11), (II.12) and (II.13), we arrive at

$$\begin{aligned}\frac{1}{\pi}\left(\frac{\|a - \hat{a}/\|\hat{a}\|_2\|_2}{\langle a, \hat{a}/\|\hat{a}\|_2\rangle}\frac{1}{2}e^{-1/2} - \frac{\sqrt{2\pi}}{r}\right) \\ \leq \mathbb{E}(f_a(\hat{a}) - f_a(a)) \leq \varepsilon,\end{aligned}$$

which finishes the proof of the theorem. ∎

III. Proofs

The main aim of this section is to prove Theorems II.1 and II.2. Before we come to that, we shall give a number of helpful Lemmas.

A. Concentration of $f_a(w)$

In this subsection we want to show that $f_a(w)$ does not deviate uniformly far from its expected value $\mathbb{E}f_a(w)$, i.e. we want to show that

$$\sup_{w \in K} |f_a(w) - \mathbb{E}f_a(w)|$$

is small with high probability. Therefore we will first estimate its mean

$$\mu := \mathbb{E}\left(\sup_{w \in K} |f_a(w) - \mathbb{E}f_a(w)|\right) \tag{III.1}$$

and then use a concentration inequality to prove Theorem II.1. The proof relies on standard techniques from [21] and [20] and is inspired by the analysis of 1-bit compressed sensing given in [24].

For $i = 1, \ldots, m$ let $\varepsilon_i \in \{+1, -1\}$ be i.i.d. Bernoulli variables with

$$\mathbb{P}(\varepsilon_i = 1) = \mathbb{P}(\varepsilon_i = -1) = 1/2. \tag{III.2}$$

Let us put

$$\mathcal{A}_i(w) = [1 - y_i\langle x_i, w\rangle]_+, \quad \mathcal{A}(w) = [1 - y\langle x, w\rangle]_+, \tag{III.3}$$

where x is an independent copy of any of the x_i and $y = \operatorname{sign}(\langle x, a\rangle)$. Further, we will make use of the following lemmas.

Lemma III.1. *For $m \in \mathbb{N}$, i.i.d. Bernoulli variables $\varepsilon_1, \ldots, \varepsilon_m$ according to* (III.2) *and any scalars $\lambda_1, \ldots, \lambda_m \in \mathbb{R}$ it holds*

$$\mathbb{P}\left(\sum_{i=1}^m \varepsilon_i[\lambda_i]_+ \geq t\right) \leq 2\mathbb{P}\left(\sum_{i=1}^m \varepsilon_i\lambda_i \geq t\right). \tag{III.4}$$

Proof: First we observe

$$\begin{aligned}\mathbb{P}\left(\sum_{i=1}^m \varepsilon_i[\lambda_i]_+ \geq t\right) &= \mathbb{P}\left(\sum_{\lambda_i \geq 0} \varepsilon_i\lambda_i \geq t\right) \\ &= \mathbb{P}\left(\sum_{\lambda_i \geq 0} \varepsilon_i\lambda_i \geq t \text{ and } \sum_{\lambda_i < 0} \varepsilon_i\lambda_i \geq 0\right) \\ &\quad + \mathbb{P}\left(\sum_{\lambda_i \geq 0} \varepsilon_i\lambda_i \geq t \text{ and } \sum_{\lambda_i < 0} \varepsilon_i\lambda_i < 0\right).\end{aligned}$$

Now we can estimate the second of these two probabilities by the first one and we arrive at

$$\mathbb{P}\Big(\sum_{i=1}^{m}\varepsilon_i[\lambda_i]_+ \geq t\Big) \leq 2\mathbb{P}\Big(\sum_{\lambda_i\geq 0}\varepsilon_i\lambda_i \geq t \text{ and } \sum_{\lambda_i<0}\varepsilon_i\lambda_i \geq 0\Big)$$
$$\leq 2\mathbb{P}\Big(\sum_{i=1}^{m}\varepsilon_i\lambda_i \geq t\Big).$$

■

Lemma III.2. 1) *For Gaussian random variables $x_1,\dots,x_m \in \mathbb{R}^d$ according to* (II.3) *it holds*

$$\mathbb{E}\left\|\frac{1}{m}\sum_{i=1}^{m}x_i\right\|_\infty \leq \frac{r\sqrt{2\log(2d)}}{\sqrt{m}}. \tag{III.5}$$

2) *Let the i.i.d. Bernoulli variables $\varepsilon_1,\dots,\varepsilon_m$ be according to* (III.2) *and let $u>0$. Then it holds*

$$\mathbb{P}\Big(\Big|\frac{1}{m}\sum_{i=1}^{m}\varepsilon_i\Big| \geq u\Big) \leq 2\exp\Big(\frac{-mu^2}{2}\Big). \tag{III.6}$$

3) *For $x_1,\dots,x_m \in \mathbb{R}^d$ and $K\subset\mathbb{R}^d$ according to* (II.3) *and* (II.6) *we denote*

$$\tilde{\mu} = \mathbb{E}\Big(\sup_{w\in K}\Big\langle\frac{1}{m}\sum_{i=1}^{m}x_i, w\Big\rangle\Big). \tag{III.7}$$

Then it holds

$$\mathbb{P}\Big(\sup_{w\in K}\Big\langle\frac{1}{m}\sum_{i=1}^{m}x_i, w\Big\rangle \geq \tilde{\mu}+u\Big) \leq \exp\Big(\frac{-mu^2}{2r^2R^2}\Big). \tag{III.8}$$

Proof:

1) The statement follows from

$$\mathbb{E}\left\|\frac{1}{m}\sum_{i=1}^{m}x_i\right\|_\infty = \frac{r}{\sqrt{m}}\mathbb{E}\|\tilde{x}\|_\infty$$

with $\tilde{x}\sim\mathcal{N}(0,\mathrm{Id})$ and proposition 8.1 of [13]:

$$\frac{\sqrt{\log(d)}}{4} \leq \mathbb{E}\|\tilde{x}\|_\infty \leq \sqrt{2\log(2d)}. \tag{III.9}$$

2) The estimate follows as a consequence of Hoeffding's inequality [16].

3) Theorem 5.2 of [24] gives the estimate

$$\mathbb{P}\Big(\sup_{w\in K}\Big\langle\frac{1}{m}\sum_{i=1}^{m}x_i, w\Big\rangle \geq \tilde{\mu}+u\Big) \leq \exp\Big(\frac{-u^2}{2\sigma^2}\Big)$$

with

$$\sigma^2 = \sup_{w\in K}\mathbb{E}\Big(\Big\langle\frac{1}{m}\sum_{i=1}^{m}x_i, w\Big\rangle^2\Big).$$

Since the $x_i's$ are independent we get

$$\frac{1}{m}\sum_{i=1}^{m}x_i = \frac{r}{\sqrt{m}}\tilde{x} \quad \text{with} \quad \tilde{x}\sim\mathcal{N}(0,\mathrm{Id})$$

and we end up with

$$\sigma^2 = \sup_{w\in K}\mathbb{E}\Big(\frac{r^2}{m}\langle\tilde{x},w\rangle^2\Big) = \frac{r^2}{m}\sup_{w\in K}\|w\|_2^2 = \frac{r^2R^2}{m}. \tag{III.10}$$

■

1) Estimate of the mean μ: To estimate the mean μ, we first derive the following symmetrization inequality, cf. [21, Chapter 6] and [24, Lemma 5.1].

Lemma III.3 (Symmetrization)**.** *Let $\varepsilon_1, \dots, \varepsilon_m$ be i.i.d. Bernoulli variables according to* (III.2)*. Under the "Standing assumptions" it holds for μ defined by* (III.1)

$$\mu \le 2\mathbb{E} \sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \varepsilon_i [1 - y_i \langle x_i, w\rangle]_+ \right|. \tag{III.11}$$

Proof: Let $\mathcal{A}_i(w)$ and $\mathcal{A}(w)$ be according to (III.3). Let x_i' and x' be independent copies of x_i and x. Then $\mathcal{A}_i'(w)$ and $\mathcal{A}'(w)$, generated in the same way (III.3) with x_i' and x' instead of x_i and x, are independent copies of $\mathcal{A}_i(w)$ and $\mathcal{A}(w)$. We denote by $\mathbb{E}'$ the mean value with respect to x_i' and x'. Using $\mathbb{E}'\big(\mathcal{A}_i'(w) - \mathbb{E}'\mathcal{A}'(w)\big) = 0$, we get

$$\begin{aligned}
\mu &= \mathbb{E} \sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \big(\mathcal{A}_i(w) - \mathbb{E}\mathcal{A}(w)\big) \right| \\
&= \mathbb{E} \sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \big(\mathcal{A}_i(w) - \mathbb{E}\mathcal{A}(w)\big) \right. \\
&\qquad\qquad \left. - \mathbb{E}'\big(\mathcal{A}_i'(w) - \mathbb{E}'\mathcal{A}'(w)\big) \right| \\
&= \mathbb{E} \sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \mathbb{E}'\big(\mathcal{A}_i(w) - \mathcal{A}_i'(w)\big) \right|.
\end{aligned}$$

Applying Jensen's inequality we further get

$$\begin{aligned}
\mu &\le \mathbb{E}\,\mathbb{E}' \sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \big(\mathcal{A}_i(w) - \mathcal{A}_i'(w)\big) \right| \\
&= \mathbb{E}\,\mathbb{E}' \sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \varepsilon_i \big(\mathcal{A}_i(w) - \mathcal{A}_i'(w)\big) \right| \\
&\le 2\mathbb{E} \sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \varepsilon_i \mathcal{A}_i(w) \right| \\
&= 2\mathbb{E} \sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \varepsilon_i [1 - y_i \langle x_i, w\rangle]_+ \right|
\end{aligned}$$

as claimed. ∎

Equipped with this tool, we deduce the following estimate for μ.

Lemma III.4. *Under the "Standing assumptions" we have*

$$\mu = \mathbb{E} \sup_{w\in K} |f_a(w) - \mathbb{E} f_a(w)| \le \frac{4\sqrt{8\pi} + 8rR\sqrt{2\log(2d)}}{\sqrt{m}}.$$

Proof: Using Lemma III.3 we obtain

$$\begin{aligned}
\mu &= \mathbb{E} \sup_{w\in K} |f_a(w) - \mathbb{E} f_a(w)| \\
&\le 2\mathbb{E} \sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \varepsilon_i [1 - y_i \langle x_i, w\rangle]_+ \right| \\
&= 2 \int_0^\infty \mathbb{P}\left(\sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \varepsilon_i [1 - y_i \langle x_i, w\rangle]_+ \right| \ge t \right) dt.
\end{aligned}$$

Now we can apply Lemma III.1 to get

$$\begin{aligned}
\mu &\le 4 \int_0^\infty \mathbb{P}\left(\sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \varepsilon_i (1 - y_i \langle x_i, w\rangle) \right| \ge t \right) dt \\
&\le 4 \int_0^\infty \mathbb{P}\left(\left| \frac{1}{m} \sum_{i=1}^{m} \varepsilon_i \right| \ge t/2 \right) \\
&\qquad + \mathbb{P}\left(\sup_{w\in K} \left| \frac{1}{m} \sum_{i=1}^{m} \varepsilon_i y_i \langle x_i, w\rangle \right| \ge t/2 \right) dt.
\end{aligned}$$

Using the second part of Lemma III.2 we can further estimate

$$\mu \leq \frac{4\sqrt{8\pi}}{\sqrt{m}} + 4\int_0^\infty \mathbb{P}\bigg(\sup_{w\in K}\Big|\frac{1}{m}\sum_{i=1}^m \varepsilon_i y_i\langle x_i, w\rangle\Big| \geq t/2\bigg)\,dt$$
$$= \frac{4\sqrt{8\pi}}{\sqrt{m}} + 8\mathbb{E}\bigg(\sup_{w\in K}\Big|\Big\langle \frac{1}{m}\sum_{i=1}^m \varepsilon_i x_i, w\Big\rangle\Big|\bigg).$$

Using the duality $\|\cdot\|_1' = \|\cdot\|_\infty$ and the first part of Lemma III.2 we get

$$= \frac{4\sqrt{8\pi}}{\sqrt{m}} + 8R\,\mathbb{E}\bigg\|\frac{1}{m}\sum_{i=1}^m x_i\bigg\|_\infty \leq \frac{4\sqrt{8\pi}}{\sqrt{m}} + \frac{8rR\sqrt{2\log(2d)}}{\sqrt{m}}.$$

■

2) Concentration inequalities: In this subsection we will estimate the probability that $f_a(w)$ deviates anywhere on K far from its mean, i.e. the probability

$$\mathbb{P}\bigg(\sup_{w\in K}|f_a(w) - \mathbb{E}f_a(w)| \geq \mu + t\bigg)$$

for some $t > 0$. First we obtain the following modified version of the second part of Lemma 5.1 of [24], cf. also [21, Chapter 6.1].

Lemma III.5 (Deviation inequality)**.** *Let $\varepsilon_1, \ldots, \varepsilon_m$ be i.i.d. Bernoulli variables according to* (III.2) *and let the "Standing assumptions" be fulfilled. Then, for $\mu \in \mathbb{R}$ according to* (III.1) *and any $t > 0$, it holds*

$$\mathbb{P}\bigg(\sup_{w\in K}|f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu + t\bigg) \tag{III.12}$$
$$\leq 4\mathbb{P}\bigg(\sup_{w\in K}\Big|\frac{1}{m}\sum_{i=1}^m \varepsilon_i[1 - y_i\langle x_i, w\rangle]_+\Big| \geq t/2\bigg).$$

Proof: Using Markov's inequality let us first note

$$\mathbb{P}\bigg(\sup_{w\in K}|f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu\bigg)$$
$$\leq \frac{\mathbb{E}\sup_{w\in K}|f_a(w) - \mathbb{E}f_a(w)|}{2\mu} = \frac{1}{2}.$$

Using this inequality we get

$$\frac{1}{2}\mathbb{P}\bigg(\sup_{w\in K}|f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu + t\bigg)$$
$$\leq \bigg(1 - \mathbb{P}\bigg(\sup_{w\in K}|f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu\bigg)\bigg)$$
$$\cdot\,\mathbb{P}\bigg(\sup_{w\in K}|f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu + t\bigg)$$
$$= \mathbb{P}\bigg(\forall w \in K : |f_a(w) - \mathbb{E}f_a(w)| < 2\mu\bigg)$$
$$\cdot\,\mathbb{P}\bigg(\exists w \in K : |f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu + t\bigg).$$

Let $\mathcal{A}_i$ and ε_i be again defined by (III.2), (III.3) and let $\mathcal{A}_i'$ be independent copies of $\mathcal{A}_i$. We further get

$$\frac{1}{2}\mathbb{P}\bigg(\sup_{w\in K}|f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu + t\bigg)$$
$$\leq \mathbb{P}\bigg(\forall w \in K : \Big|\frac{1}{m}\sum_{i=1}^m\Big(\mathcal{A}_i(w) - \mathbb{E}\mathcal{A}(w)\Big)\Big| < 2\mu\bigg)$$
$$\cdot\,\mathbb{P}\bigg(\exists w \in K : \Big|\frac{1}{m}\sum_{i=1}^m\Big(\mathcal{A}_i'(w) - \mathbb{E}\mathcal{A}'(w)\Big)\Big| \geq 2\mu + t\bigg)$$
$$\leq \mathbb{P}\bigg(\exists w \in K : \Big|\frac{1}{m}\sum_{i=1}^m\Big(\big(\mathcal{A}_i(w) - \mathbb{E}\mathcal{A}(w)\big)$$

$$
\begin{aligned}
&\quad - \big(\mathcal{A}_i'(w) - \mathbb{E}\mathcal{A}'(w)\big)\Big)\Big| \geq t\Big) \\
&= \mathbb{P}\Big(\exists w \in K : \Big|\frac{1}{m}\sum_{i=1}^m \varepsilon_i(\mathcal{A}_i(w) - \mathcal{A}_i'(w))\Big| \geq t\Big) \\
&\leq 2\mathbb{P}\Big(\exists w \in K : \Big|\frac{1}{m}\sum_{i=1}^m \varepsilon_i \mathcal{A}_i(w)\Big| \geq t/2\Big),
\end{aligned}
$$

which yields the claim. ∎

Combining the Lemmas III.1 and III.5 we deduce the following result.

Lemma III.6. *Under the "Standing assumptions" it holds for μ and $\tilde{\mu}$ according to* (III.1) *and* (III.7) *and any $u > 0$*

$$
\begin{aligned}
&\mathbb{P}\Big(\sup_{w\in K} |f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu + 2\tilde{\mu} + u\Big) \\
&\quad \leq 8\Big(\exp\Big(\frac{-mu^2}{32}\Big) + \exp\Big(\frac{-mu^2}{32r^2R^2}\Big)\Big).
\end{aligned} \tag{III.13}
$$

Proof: Applying Lemma III.5 and Lemma III.1 we get

$$
\begin{aligned}
&\mathbb{P}\Big(\sup_{w\in K} |f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu + 2\tilde{\mu} + u\Big) \\
&\leq 4\mathbb{P}\Big(\sup_{w\in K}\Big|\frac{1}{m}\sum_{i=1}^m \varepsilon_i[1 - y_i\langle x_i, w\rangle]_+\Big| \geq \tilde{\mu} + u/2\Big) \\
&\leq 8\mathbb{P}\Big(\sup_{w\in K}\Big|\frac{1}{m}\sum_{i=1}^m \varepsilon_i(1 - y_i\langle x_i, w\rangle)\Big| \geq \tilde{\mu} + u/2\Big) \\
&\leq 8\mathbb{P}\Big(\Big|\frac{1}{m}\sum_{i=1}^m \varepsilon_i\Big| \geq u/4\Big) \\
&\qquad + 8\mathbb{P}\Big(\sup_{w\in K}\Big|\Big\langle \frac{1}{m}\sum_{i=1}^m x_i, w\Big\rangle\Big| \geq \tilde{\mu} + u/4\Big).
\end{aligned}
$$

Finally, applying the second and third part of Lemma III.2 this can be further estimated from above by

$$
\leq 8\Big(\exp\Big(\frac{-mu^2}{32}\Big) + \exp\Big(\frac{-mu^2}{32r^2R^2}\Big)\Big),
$$

which finishes the proof. ∎

Using the two Lemmas III.4 and III.6 we can now prove Theorem II.1.

Proof of Theorem II.1: Lemma III.6 yields

$$
\begin{aligned}
&\mathbb{P}\Big(\sup_{w\in K} |f_a(w) - \mathbb{E}f_a(w)| \geq 2\mu + 2\tilde{\mu} + u\Big) \\
&\quad \leq 8\Big(\exp\Big(\frac{-mu^2}{32}\Big) + \exp\Big(\frac{-mu^2}{32r^2R^2}\Big)\Big).
\end{aligned}
$$

Using Lemma III.4 we further get

$$
\mu \leq \frac{4\sqrt{8\pi} + 8rR\sqrt{2\log(2d)}}{\sqrt{m}}.
$$

Invoking the duality $\|\cdot\|_1' = \|\cdot\|_\infty$ and the first part of Lemma III.2 we can further estimate $\tilde{\mu}$ by

$$
\tilde{\mu} = R\mathbb{E}\left\|\frac{1}{m}\sum_{i=1}^m x_i\right\|_\infty \leq \frac{rR\sqrt{2\log(2d)}}{\sqrt{m}}.
$$

Hence, with probability at least

$$
1 - 8\Big(\exp\Big(\frac{-mu^2}{32}\Big) + \exp\Big(\frac{-mu^2}{32r^2R^2}\Big)\Big)
$$

we have

$$\sup_{w\in K}|f_a(w)-\mathbb{E}f_a(w)| \le 2\mu+2\tilde{\mu}+u$$
$$\le \frac{8\sqrt{8\pi}+18rR\sqrt{2\log(2d)}}{\sqrt{m}}+u$$

as claimed. ∎

B. Estimate of the expected value

In this subsection we will estimate

$$\mathbb{E}(f_a(w)-f_a(a)) = \mathbb{E}[1-y\langle x,w\rangle]_+ - \mathbb{E}[1-y\langle x,a\rangle]_+$$

for some $w\in\mathbb{R}^d\backslash\{0\}$ with $\|w\|_1\le R$. We will first calculate both expected values separately and later estimate their difference. We will make use of the following statements from probability theory.

Lemma III.7. *Let $a,x\in\mathbb{R}^d$ be according to* (II.4)*,* (II.3) *and let $w\in\mathbb{R}^d\backslash\{0\}$. Then it holds*

1) $\langle x,a\rangle,\ \langle x,\frac{w}{\|w\|_2}\rangle\sim\mathcal{N}(0,r^2)$,
2) $\mathrm{Cov}(\langle x,a\rangle,\langle x,w\rangle)=r^2\langle a,w\rangle$.

Proof: The first statement is well known in probability theory as the 2-stability of normal distribution. For the second statement we get

$$\mathrm{Cov}(\langle x,a\rangle,\langle x,w\rangle)=\mathbb{E}(\langle x,a\rangle\langle x,w\rangle)=\sum_{i,j=1}^{d}a_iw_j\mathbb{E}(x_ix_j)$$
$$=r^2\sum_{i=1}^{d}a_iw_i=r^2\langle a,w\rangle$$

as claimed. ∎

It is very well known, cf. [14, Corollary 5.2], that projections of a Gaussian random vector onto two orthogonal directions are mutually independent.

Lemma III.8. *Let $x\sim\mathcal{N}(0,\mathrm{Id})$ and let $a,b\in\mathbb{R}^d$ with $\langle a,b\rangle=0$. Then $\langle x,a\rangle$ and $\langle x,b\rangle$ are independent random variables.*

Applying these two lemmas to our case we end up with the following lemma.

Lemma III.9. *For $a\in\mathbb{R}^d$ according to* (II.4)*, $x\sim\mathcal{N}(0,r^2\mathrm{Id})$ and $w\in\mathbb{R}^d$ we have*

$$\langle x,w\rangle = c\langle x,a\rangle + c'Z$$

for some $Z\sim\mathcal{N}(0,r^2)$ independent of $\langle x,a\rangle$ and

$$c=\langle a,w\rangle,\quad c'=\sqrt{\|w\|_2^2-c^2}. \tag{III.14}$$

Remark III.10. Note that c' is well defined, since $c^2\le\|w\|_2^2\|a\|_2^2=\|w\|_2^2$.

Proof: If $c'=0$, the statement holds trivially. If $c'\neq0$, we set

$$Z=\frac{1}{c'}(\langle x,w\rangle-c\langle x,a\rangle)=\frac{1}{c'}\sum_{i=1}^{d}x_i\left(w_i-ca_i\right).$$

Hence, Z is indeed normally distributed with $\mathbb{E}(Z)=0$ and $\mathrm{Var}(Z)=r^2$. It remains to show that Z and $\langle x,a\rangle$ are independent. We observe that

$$\langle a,w-ca\rangle=\langle a,w\rangle-\langle a,w\rangle\|a\|_2=0$$

and, finally, Lemma III.8 yields the claim. ∎

Lemma III.11. *Let $a\in\mathbb{R}^d$ and $f_a\colon\mathbb{R}^d\to\mathbb{R}$ be according to* (II.4)*,* (II.7)*. Then it holds*

1) $\mathbb{E}f_a(a)=\frac{1}{\sqrt{2\pi}}\int_{\mathbb{R}}\left[1-r|t|\right]_+e^{\frac{-t^2}{2}}\,dt$,
2) $\mathbb{E}f_a(w)=\frac{1}{2\pi}\int_{\mathbb{R}^2}\left[1-cr|t_1|-c'rt_2\right]_+e^{\frac{-t_1^2-t_2^2}{2}}\,dt_1\,dt_2$, *where c and c' are defined by* (III.14).

Proof:

1) Let $\omega \sim \mathcal{N}(0,1)$ and use the first part of Lemma III.7 to obtain

$$\begin{aligned}\mathbb{E}f_a(a) &= \mathbb{E}[1-|\langle x,a\rangle|]_+ = \mathbb{E}[1-r|\omega|]_+\\ &= \frac{1}{\sqrt{2\pi}}\int_{\mathbb{R}}\left[1-r|t|\right]_+ e^{\frac{-t^2}{2}}\,dt.\end{aligned}$$

2) Using the notation of Lemma III.9 we get

$$\begin{aligned}\mathbb{E}f_a(w) &= \mathbb{E}[1-\operatorname{sign}(\langle x,a\rangle)\langle x,w\rangle]_+\\ &= \mathbb{E}[1-\operatorname{sign}(\langle x,a\rangle)(c\langle x,a\rangle + c'Z)]_+\\ &= \mathbb{E}[1-c|\langle x,a\rangle| - c'\operatorname{sign}(\langle x,a\rangle)Z]_+\\ &= \mathbb{E}[1-c|\langle x,a\rangle| - c'Z]_+\\ &= \frac{1}{2\pi}\int_{\mathbb{R}^2}[1-cr|t_1|-c'rt_2]_+ e^{\frac{-t_1^2-t_2^2}{2}}\,dt_1\,dt_2.\end{aligned}$$

■

Using this result we now can prove Theorem II.2.

Proof of Theorem II.2: Using Lemma III.11 we first observe

$$\begin{aligned}-\pi\mathbb{E}f_a(a) &= -\frac{\sqrt{\pi}}{\sqrt{2}}\int_{\mathbb{R}}[1-r|t|]_+ e^{\frac{-t^2}{2}}\,dt \qquad \text{(III.15)}\\ &= -\sqrt{2\pi}\int_0^{\frac{1}{r}}\left(1-rt\right)e^{\frac{-t^2}{2}}\,dt\\ &\geq -\sqrt{2\pi}\int_0^{\frac{1}{r}} e^{\frac{-t^2}{2}}\,dt \geq -\frac{\sqrt{2\pi}}{r}.\end{aligned}$$

To estimate the expected value of $f_a(w)$ we now distinguish the two cases $c \leq 0$ and $c > 0$.
1. case: $c \leq 0$*:* In that case we get

$$\pi\mathbb{E}f_a(w) = \int_{\mathbb{R}}\int_0^{\infty}\left[1-crt_1-c'rt_2\right]_+ e^{\frac{-t_1^2-t_2^2}{2}}\,dt_1\,dt_2.$$

Since $-crt_1 \geq 0$ for $0 \leq t_1 < \infty$ we can further estimate

$$\begin{aligned}\pi\mathbb{E}f_a(w) &\geq \int_{\mathbb{R}}\int_0^{\infty}\left[1-c'rt_2\right]_+ e^{\frac{-t_1^2-t_2^2}{2}}\,dt_1\,dt_2\\ &\geq \int_{-\infty}^{0}\int_0^{\infty}(1-c'rt_2)e^{\frac{-t_1^2-t_2^2}{2}}\,dt_1\,dt_2\\ &= \int_{-\infty}^{0}\int_0^{\infty} e^{\frac{-t_1^2-t_2^2}{2}}\,dt_1\,dt_2 + c'r\int_0^{\infty}\int_0^{\infty} t_2 e^{\frac{-t_1^2-t_2^2}{2}}\,dt_1\,dt_2\\ &= \frac{\pi}{2} + c'r\frac{\sqrt{\pi}}{\sqrt{2}}.\end{aligned}$$

As claimed, putting both terms together, we arrive at

$$\pi\mathbb{E}(f_a(w)-f_a(a)) \geq \frac{\pi}{2} + c'r\frac{\sqrt{\pi}}{\sqrt{2}} - \frac{\sqrt{2\pi}}{r}.$$

2. case: $c > 0$*:* First let us observe that $1-crt_1-c'rt_2 \geq 0$ on $[0,1/cr]\times(-\infty,0] \subset \mathbb{R}^2$. Hence, we get

$$\begin{aligned}\pi\mathbb{E}f_a(w) &= \int_{\mathbb{R}^2}[1-crt_1-c'rt_2]_+ e^{\frac{-t_1^2-t_2^2}{2}}\,dt_2\,dt_1\\ &\geq \int_0^{\frac{1}{cr}}\int_{-\infty}^{0}(1-crt_1-c'rt_2)e^{\frac{-t_1^2-t_2^2}{2}}\,dt_2\,dt_1\\ &= \frac{\sqrt{\pi}}{\sqrt{2}}\int_0^{\frac{1}{cr}}(1-crt)e^{\frac{-t^2}{2}}\,dt + c'r\int_0^{\frac{1}{cr}} e^{\frac{-t^2}{2}}\,dt\\ &\geq \frac{\sqrt{\pi}}{\sqrt{2}}\int_0^{\frac{1}{cr}}(1-crt)e^{\frac{-t^2}{2}}\,dt + \frac{c'}{c}\exp\left(\frac{-1}{2c^2r^2}\right).\end{aligned}$$

Combining this estimate with (III.15) we arrive at

$$\pi\mathbb{E}(f_a(w)-f_a(a)) \geq \frac{\sqrt{\pi}}{\sqrt{2}}\int_0^{\frac{1}{cr}}(1-crt)e^{\frac{-t^2}{2}}\,dt + \frac{c'}{c}\exp\left(\frac{-1}{2c^2r^2}\right)-\frac{\sqrt{2\pi}}{r}.$$

■

IV. ℓ_1-SVM WITH ADDITIONAL ℓ_2-CONSTRAINT

A detailed inspection of the analysis done so far shows that it would be convenient if the convex body K would not include vectors with large ℓ_2-norm. For example, in (III.10) we needed to calculate $\sup_{w\in K}\|w\|_2^2 = R^2$, although the measure of the set of vectors in K with ℓ_2-norm close to R is extremely small.

Therefore, we will modify the ℓ_1-SVM (II.5) by adding an additional ℓ_2-constraint, that is instead of (II.5) we consider the optimization problem

$$\min_{w\in\mathbb{R}^d}\sum_{i=1}^m[1-y_i\langle x_i,w\rangle]_+ \text{ s. t. } \|w\|_1\leq R \text{ and } \|w\|_2\leq 1. \tag{IV.1}$$

The combination of ℓ_1 and ℓ_2 constraints is by no means new - for example, it plays a crucial role in the theory of elastic nets [32]. Furthermore, let us remark that the set

$$\tilde{K} = \{w\in\mathbb{R}^d \mid \|w\|_1\leq R \text{ and } \|w\|_2\leq 1\} \tag{IV.2}$$

appears also in [24]. We get $\tilde{K}\subset K$ with K according to (II.6). Hence, Theorem II.1 and (II.8) still remain true if we replace K by $\tilde{K}$ and we obtain

$$\sup_{w\in\tilde{K}}|f_a(w)-\mathbb{E}f_a(w)| \leq \frac{8\sqrt{8\pi}+18rR\sqrt{2\log(2d)}}{\sqrt{m}}+u \tag{IV.3}$$

with high probability and

$$\mathbb{E}(f_a(\hat{a})-f_a(a)) \leq 2\sup_{w\in\tilde{K}}|f_a(\hat{a})-f_a(a)|, \tag{IV.4}$$

where $\hat{a}$ is now the minimizer of (IV.1).

It remains to estimate the expected value $\mathbb{E}(f_a(w)-f_a(a))$ in order to obtain an analogue of Theorem II.3 for (IV.1), which reads as follows.

Theorem IV.1. *Let $d\geq 2$, $0<\varepsilon<1/2$, $r>2\sqrt{2\pi}(1-2\varepsilon)^{-1}$, $a\in\mathbb{R}^d$ according to* (II.4)*, $m\geq C\varepsilon^{-2}r^2R^2\log(d)$ for some constant C, $x_1,\dots,x_m\in\mathbb{R}^d$ according to* (II.3) *and $\hat{a}\in\mathbb{R}^d$ a minimizer of* (IV.1)*. Then it holds*

$$\|a-\hat{a}\|_2^2 \leq \frac{C'\varepsilon}{r(1-\exp\left(\frac{-1}{2r^2}\right))} \tag{IV.5}$$

with probability at least

$$1-\gamma\exp\left(-C''\log(d)\right)$$

for some positive constants γ, C', C''.

Remark IV.2. 1) As for Theorem II.3 we can write down the expressions explicitly, i.e. without the constants γ, C, C' and C''. That is, taking $m\geq 4\varepsilon^{-2}\left(8\sqrt{8\pi}+(18+t)rR\sqrt{2\log(2d)}\right)^2$ for some $t>0$, we get

$$\|a-\hat{a}\|_2^2 \leq \frac{\sqrt{\pi/2}\,\varepsilon}{r\left(1-\exp\left(\frac{-1}{2r^2}\right)\right)}.$$

with probability at least

$$1-8\left(\exp\left(\frac{-t^2r^2R^2\log(2d)}{16}\right)+\exp\left(\frac{-t^2\log(2d)}{16}\right)\right).$$

2) The main advantage of Theorem IV.1 compared to Theorem II.3 is that the parameter r does not need to grow to infinity. Actually, (IV.5) is clearly not optimal for large r. Indeed, if (say) $\varepsilon<0.2$, we can take $r=10$, and obtain

$$\|a-\hat{a}\|_2^2 \leq \tilde{C}'\varepsilon$$

for $m \geq \tilde{C}\varepsilon^{-2}R^2\log(d)$ with high probability.

Proof: As in the proof of Theorem II.3 we first obtain $c' = \sqrt{\|\hat{a}\|_2^2 - \langle a, \hat{a}\rangle^2} > 0$ and $c = \langle a, \hat{a}\rangle > 0$. Using Lemma III.11 we get

$$\begin{aligned}
&\pi\mathbb{E}(f_a(w) - f_a(a)) \\
&\geq \int_0^{\frac{1}{r}} \int_{\mathbb{R}} \big((1 - crt_1 - c'rt_2) - (1 - rt_1)\big) e^{\frac{-t_1^2 - t_2^2}{2}}\, dt_2\, dt_1 \\
&= r(1-c)\sqrt{2\pi} \int_0^{\frac{1}{r}} t e^{\frac{-t^2}{2}}\, dt
\end{aligned}$$

with

$$1 - c = 1 - \langle a, \hat{a}\rangle \geq \frac{1}{2}(\|a\|_2^2 + \|\hat{a}\|_2^2) - \langle a, \hat{a}\rangle = \frac{1}{2}\|a - \hat{a}\|_2^2.$$

The claim now follows from (IV.4) and (IV.3). ∎

V. Numerical experiments

We performed several numerical tests to exhibit different aspects of the algorithms discussed above. In the first two parts of this section we fixed $d = 1000$ and set $\tilde{a} \in \mathbb{R}^d$ with 5 nonzero entries $\tilde{a}_{10} = 1$, $\tilde{a}_{140} = -1$, $\tilde{a}_{234} = 0.5$, $\tilde{a}_{360} = -0.5$, $\tilde{a}_{780} = 0.3$, Afterwards we normalized $\tilde{a}$ and set $a = \tilde{a}/\|\tilde{a}\|_2$ and $R = \|a\|_1$.

A. Dependency on r

We run the ℓ_1-SVM (II.5) with $m = 200$ and $m = 400$ for different values of r between zero and 1.5. The same was done for the ℓ_1-SVM with the additional ℓ_2-constraint (IV.1), which is called $\ell_{1,2}$-SVM in the legend of the figure. The average error of $n = 20$ trials between a and $\hat{a}/\|\hat{a}\|_2$ is plotted against r. We observe that especially for small r's the ℓ_1-SVM with ℓ_2-constraint performs much better than classical ℓ_1-SVM.

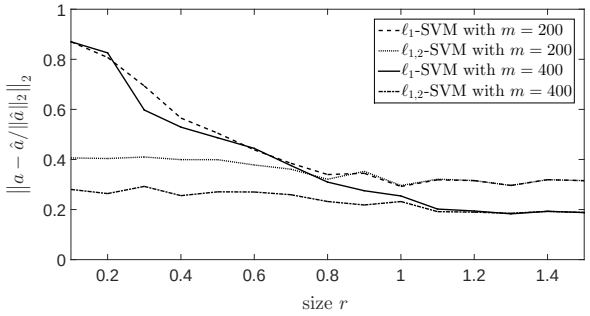

Figure 1. Dependency on r

B. Dependency on m and comparison with 1-Bit CS

In the second experiment, we run ℓ_1-SVM with and without the extra ℓ_2-constraint for two different values of r, namely for $r = 0.75$ and for r depending on m as $r = \sqrt{m}/30$. We plotted the average error of $n = 40$ trials for each value. The last method used is 1-bit Compressed Sensing [24], which is given as the maximizer of

$$\max_{w \in \mathbb{R}^d} \sum_{i=1}^m y_i \langle x_i, w\rangle \quad \text{subject to} \quad \|w\|_2 \leq 1,\ \|w\|_1 \leq R. \tag{V.1}$$

Note that maximizer of (V.1) is independent of r, since it is linear in x_i. First, one observes that the error of ℓ_1-SVM does not converge to zero if the value of $r = 0.75$ is fixed. This is in a good agreement with Theorem II.3 and the error estimate (II.9). This drawback disappears when $r = \sqrt{m}/30$ grows with m, but ℓ_1-SVM still performs quite badly. The two versions of $\ell_{1,2}$-SVM perform essentially better than ℓ_1-SVM, and slightly better than 1-bit Compressed Sensing.

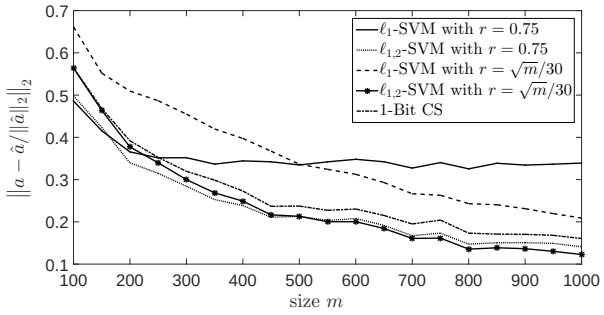

Figure 2. Comparison of ℓ_1-SVM with 1-Bit CS.

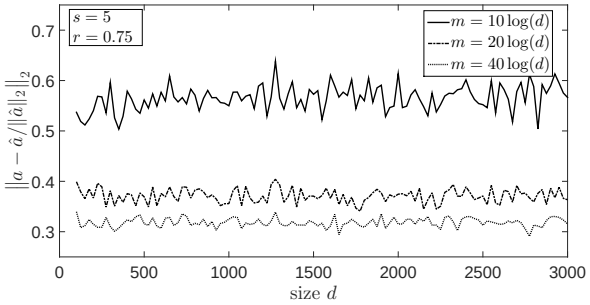

Figure 3. Dependency on d.

C. Dependency on d

In figure 3 we investigated the dependency of the error of ℓ_1-SVM on the dimension d. We fixed the sparsity level $s = 5$ and for each d between 100 and 3000 we draw an s-sparse signal a and measurement vectors x_i at random. Afterwards we run the ℓ_1-SVM with the three different values $m = m_i \log(d)$ with $m_1 = 10$, $m_2 = 20$ and $m_3 = 40$. We plotted the average errors between a and $\hat{a}/\|\hat{a}\|_2$ for $n = 60$ trials.

We indeed see that to achieve the same error, the number of measurements only needs to grow logarithmically in d, explaining once again the success of ℓ_1-SVM for high-dimensional classification problems.

VI. Discussion

In this paper we have analyzed the performance of ℓ_1-SVM (II.5) in recovering sparse classifiers. Theorem II.3 shows, that a good approximation of such a sparse classifier can be achieved with small number of learning points m if the data is well spread. The geometric properties of well distributed learning points are modelled by independent Gaussian vectors with growing variance r and it would be interesting to know, how ℓ_1-SVM performs on points chosen independently from other distributions. The number of learning points needs to grow logarithmically with the underlying dimension d and linearly with the sparsity of the classifier. On the other hand, the optimality of the dependence of m on ε and r remains open. Another important question left open is the behavior of ℓ_1-SVM in the presence of missclasifications, i.e. when there is a (small) probability that the signs $y_i \in \{-1, +1\}$ do not coincide with $\mathrm{sign}(\langle x_i, a\rangle)$. Finally, we proposed a modification of ℓ_1-SVM by incorporating an additional ℓ_2-constraint.

Acknowledgment

We would like to thank A. Hinrichs, M. Omelka, and R. Vershynin for valuable discussions.

References

[1] A. Ai, A. Lapanowski, Y. Plan, and R. Vershynin, "One-bit compressed sensing with non-Gaussian measurements", *Linear Algebra Appl.*, vol. 441, pp. 222–239, 2014.
[2] K.P. Bennett and O.L. Mangasarian, "Robust linear programming discrimination of two linearly independent inseparable sets", *Optimization Methods and Software*, pp. 23–34, 1992.
[3] H. Boche, R. Calderbank, G. Kutyniok, and J. Vybíral, "A survey of compressed sensing", Applied and Numerical Harmonic Analysis, Birkhäuser, Boston, 2015.
[4] P.T. Boufounos and R.G. Baraniuk, "1-Bit compressive sensing", In 42nd Annual Conference on Information Sciences and Systems, 2008.
[5] P.S. Bradley and O.L. Mangasarian, "Feature selection via mathematical programming", *INFORMS J. Comput.*, vol. 10, pp. 209–217, 1998.
[6] P.S. Bradley and O.L. Mangasarian, "Feature selection via concave minimization and support vector machines", In Proceedings of the 13th International Conference on Machine Learning, pp. 82–90, 1998.

[7] E. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information”, *IEEE Trans. Inform. Theory*, vol. 52, pp. 489–509, 2006.
[8] C. Cortes and V. Vapnik, “Support-vector networks”, *Machine Learning*, vol. 20, no.3, pp. 273–297, 1995.
[9] F. Cucker and D. X. Zhou, *Learning Theory: An Approximation Theory Viewpoint*, Cambridge University Press, 2007.
[10] M.A. Davenport, M.F. Duarte, Y.C. Eldar, and G. Kutyniok, *Introduction to compressed sensing*, in Compressed sensing, Cambridge Univ. Press, Cambridge, pp. 1–64, 2012.
[11] D.L. Donoho, “Compressed sensing”, *IEEE Trans. Inform. Theory*, vol. 52, pp. 1289–1306, 2006.
[12] M. Fornasier and H. Rauhut, *Compressive sensing*, In: *Handbook of Mathematical Methods in Imaging*, Springer, pp. 187–228, 2011.
[13] S. Foucart and H. Rauhut, *A mathematical introduction to compressive sensing*, Applied and Numerical Harmonic Analysis, Birkhäuser, Boston, 2013.
[14] W. Härdle and L. Simar, *Applied multivariate statistical analysis*, Springer, Berlin, 2003.
[15] M. Hilario and A. Kalousis, “Approaches to dimensionality reduction in proteomic biomarker studies”, *Brief Bioinform*, vol. 9, no. 2, pp. 102–118, 2008.
[16] W. Hoeffding, “Probability inequalities for sums of bounded random variables”, *J. Amer. Stat. Assoc.*, vol. 58, pp. 13–30, 1963.
[17] K. Knudson, R. Saab, and R. Ward, *One-bit compressive sensing with norm estimation*, preprint, available at `http://arxiv.org/abs/1404.6853`.
[18] O.L. Mangasarian, “Arbitrary-norm separating plane”, *Oper. Res. Lett.*, vol. 24, pp. 15–23, 1999.
[19] O.L. Mangasarian, “Support vector machine classification via parameterless robust linear programming”, *Optim. Methods Softw.*, vol. 20, pp. 115–125, 2005.
[20] M. Ledoux, *The Concentration of Measure Phenomenon*, Am. Math. Soc., 2001.
[21] M. Ledoux and M. Talagrand, *Probability in Banach Spaces: Isoperimetry and Processes*, Springer, Berlin, 1991.
[22] Y. Plan, R. Vershynin, and E. Yudovina, “High-dimensional estimation with geometric constraints”, available at `http://arxiv.org/abs/1404.3749`.
[23] Y. Plan and R. Vershynin, “One-bit compressed sensing by linear programming”, *Comm. Pure Appl. Math.*, vol. 66, pp. 1275–1297, 2013.
[24] Y. Plan and R. Vershynin, “Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach”, *IEEE Trans. Inform. Theory*, vol. 59, pp. 482–494, 2013.
[25] I. Steinwart and A. Christmann, *Support Vector Machines*, Springer, Berlin, 2008.
[26] R. Tibshirani, “Regression shrinkage and selection via the Lasso”, *J. Royal Stat. Soc. Ser. B*, vol. 58, no. 1, pp. 267–288, 1996.
[27] V. Vapnik and A. Chervonenkis, “A note on one class of perceptrons”, *Automation and Remote Control*, vol. 25, no. 1, 1964.
[28] V. Vapnik, *The Nature of Statistical Learning Theory*, Springer, Berlin, 1995.
[29] V. Vapnik, *Statistical Learning Theory*, Wiley, Chichester, 1998.
[30] L. Wang, J. Zhu, and H. Zou, “Hybrid huberized support vector machines for microarray classification and gene selection”, *Bioinformatics*, vol. 24, no. 3, pp. 412–419, 2008.
[31] H. H. Zhang, J. Ahn, X. Lin, and Ch. Park, “Gene selection using support vector machines with non-convex penalty”, *Bioinformatics*, vol. 22, no. 1, pp. 88–95, 2006.
[32] H. Zou and T. Hastie, “Regularization and variable selection via the elastic net”, *J. R. Stat. Soc. Ser. B Stat. Methodol.*, vol. 67, no. 2, pp. 301–320, 2005.
[33] H. Zou, T. Hastie, and R. Tibshirani, “Sparse principal component analysis”, *J. Comput. Graph. Statist.*, vol. 15, no. 2, pp. 265–286, 2006.
[34] J. Zhu, S. Rosset, T. Hastie, and R. Tibshirani, “1-norm support vector machines”, In Proc. Advances in Neural Information Processing Systems, vol. 16, pp. 49–56, 2004.

Anton Kolleck received his M.S. in Mathematics at Technical University Berlin, Germany in 2013, where he now continues as Ph.D. student. His research concentrates on sparse recovery and compressed sensing and their applications in approximation theory.

Jan Vybíral received his M.S. in Mathematics at Charles Univeristy, Prague, Czech Republic in 2002. He earned the Dr. rer. nat. degree in Mathematics at Friedrich-Schiller University, Jena, Germany in 2005. He had postdoc positions in Jena, Austrian Academy of Sciences, Austria, and Technical University Berlin, Germany. He is currently an Assistant Professor of Mathematics at Charles University. His core interests are in functional analysis with applications to sparse recovery and compressed sensing.